

On Sparsity Inducing Regularization Methods for Machine Learning

Andreas Argyriou, Luca Baldassarre, Charles A. Micchelli and Massimiliano Pontil

Dedicated to Vladimir Vapnik with esteem and gratitude for his fundamental contribution to Machine Learning.

Abstract During the past years there has been an explosion of interest in learning methods based on sparsity regularization. In this paper, we discuss a general class of such methods, in which the regularizer can be expressed as the composition of a convex function ω with a linear function. This setting includes several methods such the group Lasso, the Fused Lasso, multi-task learning and many more. We present a general approach for solving regularization problems of this kind, under the assumption that the proximity operator of the function ω is available. Furthermore, we comment on the application of this approach to support vector machines, a technique pioneered by the groundbreaking work of Vladimir Vapnik.

Andreas Argyriou

Ecole Centrale de Paris, Grande Voie des Vignes 92 295 Chatenay-Malabry, FRANCE. e-mail: andreas.argyriou@ecp.fr

Luca Baldassarre

Laboratory for Information and Inference Systems, EPFL, ELD 243, Station 11 CH-1015 Lausanne, Switzerland. e-mail: luca.baldassarre@epfl.ch

Charles A. Micchelli

Department of Mathematics and Statistics, University at Albany, Earth Science 110 Albany, NY 12222, USA. e-mail: charles_micchelli@gotmail.com

Massimiliano Pontil

Department of Computer Science, University College London, Malet Place London WC1E 6BT, UK. e-mail: m.pontil@cs.ucl.ac.uk

1 Introduction

In this paper, we address supervised learning methods which are based on the optimization problem

$$\min_{x \in \mathbb{R}^d} \{f(x) + g(x)\}, \quad (1)$$

where the function f measures the fit of a vector x (linear predictor) to available training data and g is a penalty term or regularizer which encourages certain types of solutions. Specifically, we let $f(x) = E(y, Ax)$, where $E : \mathbb{R}^s \times \mathbb{R}^s \rightarrow [0, \infty)$ is an error function, $y \in \mathbb{R}^s$ is a vector of measurements and $A \in \mathbb{R}^{s \times d}$ a matrix, whose rows are the input vectors. This class of regularization methods arise in machine learning, signal processing and statistics and have a wide range of applications.

Different choices of the error function and the penalty function correspond to specific techniques. In this paper, we are interested in solving problem (1) when f is a *strongly smooth convex* function (such as the square error $E(y, Ax) = \|y - Ax\|_2^2$) and the penalty function g is obtained as the composition of a “simple” function with a linear transformation B , that is,

$$g(x) = \omega(Bx), \quad (2)$$

where B is a prescribed $m \times d$ matrix and ω is a *nondifferentiable convex* function on \mathbb{R}^d . The class of regularizers (2) includes a variety of methods, depending on the choice of the function ω and of matrix B . Our motivation for studying this class of penalty functions arises from sparsity-inducing regularization methods which consider ω to be either the ℓ_1 norm or a mixed ℓ_1 - ℓ_p norm. When B is the identity matrix and $p = 2$, the latter case corresponds to the well-known Group Lasso method [36], for which well studied optimization techniques are available. Other choices of the matrix B give rise to different kinds of Group Lasso with overlapping groups [15, 37], which have proved to be effective in modeling structured sparse regression problems. Further examples can be obtained by considering composition with the ℓ_1 norm, for example this includes the Fused Lasso penalty function [31] and the graph prediction problem of [13].

A common approach to solve many optimization problems of the general form (1) is via proximal-gradient methods. These are first-order iterative methods, whose computational cost per iteration is comparable to gradient descent. In some problems in which g has a simple expression, proximal-gradient methods can be combined with acceleration techniques [22, 24, 32], to yield significant gains in the number of iterations required to reach a certain approximation accuracy of the minimal value. The essential step of proximal-gradient methods requires the computation of the proximity operator of function g , see Definition 1 below. In certain cases of practical importance, this operator admits a closed form, which makes proximal-gradient methods appealing to use. However, in the general case (2) the proximity operator may not be easily computable.

We describe a general technique to compute the proximity operator of the composite regularizer (2) from the solution of a fixed point problem, which depends on

the proximity operator of the function ω and the matrix B . This problem can be solved by a simple and efficient iterative scheme when the proximity operator of ω has a closed form or can be computed in a finite number of steps. When f is a strongly smooth function, the above result can be used together with Nesterov's accelerated method [22, 24] to provide an efficient first-order method for solving the optimization problem (1).

The paper is organized as follows. In Section 2, we review the notion of proximity operator, useful facts from fixed point theory and present a convergent algorithm for the solution of problem (1) when f is quadratic function and then an algorithm to solve the associated optimization problem (1). In Section 3, we discuss some examples of composite functions of the form (2) which are valuable in applications. In Section 4 we apply our observations to support vector machines and obtained new algorithms for the solution of this problem. Finally, Section 5 contains concluding remarks.

2 Fixed Point Algorithms Based on Proximity Operators

In this section, we present an optimization approach which use fixed point algorithms for nonsmooth problems of the form (1) under the assumption (2). We first recall some notation and then move on to present an approach to compute the proximity operator for composite regularizers.

2.1 Notation and Problem Formulation

We denote by $\langle \cdot, \cdot \rangle$ the Euclidean inner product on \mathbb{R}^d and let $\|\cdot\|_2$ be the induced norm. If $v : \mathbb{R} \rightarrow \mathbb{R}$, for every $x \in \mathbb{R}^d$ we denote by $v(x)$ the vector $(v(x_i))_{i=1}^d$. For every $p \geq 1$, we define the ℓ_p norm of x as $\|x\|_p = (\sum_{i=1}^d |x_i|^p)^{\frac{1}{p}}$.

As the basic building block of our method, we consider the optimization problem (1) in the special case when f is a quadratic function and the regularization term g is obtained by the composition of a convex function with a linear function. That is, we consider the problem

$$\min \left\{ \frac{1}{2} y^\top Q y - x^\top y + \omega(By) : y \in \mathbb{R}^d \right\}. \quad (3)$$

where x is a given vector in \mathbb{R}^d and Q a positive definite $d \times d$ matrix. The development of a convergent method for the solution of this problem requires the well-known concepts of proximity operator and subdifferential of a convex function. Let us now review some of salient features of these important notions which are needed for the analysis of problem (3).

The proximity operator on a Hilbert space was introduced by Moreau in [20].

Definition 1. Let ω be a real valued convex function on \mathbb{R}^d . The proximity operator of ω is defined, for every $x \in \mathbb{R}^d$ by

$$\text{prox}_\omega(x) := \operatorname{argmin} \left\{ \frac{1}{2} \|y - x\|_2^2 + \omega(y) : y \in \mathbb{R}^d \right\}. \quad (4)$$

The proximity operator is well defined, because the above minimum exists and is unique.

Recall that the subdifferential of ω at x is defined as $\partial\omega(x) = \{u : u \in \mathbb{R}^d, \langle y - x, u \rangle + \omega(x) \leq \omega(y), \forall y \in \mathbb{R}^d\}$. The subdifferential is a nonempty compact and convex set. Moreover, if ω is differentiable at x then its subdifferential at x consists only of the gradient of ω at x .

The relationship between the proximity operator and the subdifferential of ω are essential for algorithmic developments for the solution of (3), [2, 9, 19, 21]. Generally the proximity operator is difficult to compute since it is expressed as the minimum of a convex optimisation problem. However, there are some rare circumstances where it can be obtained explicitly, for examples when $\omega(x)$ is a multiple of the ℓ_1 norm of x the proximity operator relates to soft thresholding and moreover a related formula allows for the explicit identification of the proximity operator for the ℓ_2 norm, see, for example, [2, 9, 19]. Our optimisation problem (3) can be reduced to the identification of the proximity operator for the composition function $\omega \circ B$. Although the prox of ω may be readily available, it may still be a computational challenge to obtain the prox of $\omega \circ B$. We consider this essential issue in the next section.

2.2 Computation of a Generalized Proximity Operator with a Fixed Point Method

In this section we consider circumstances in which the proximity operator of ω can be explicitly computed in a finite number of steps and seek an algorithm for the solution of the optimisation problem (3).

As we shall see, the method proposed here applies for any positive definite matrix Q . This will allow us in a future publication to provide a second order method for solving (1). For the moment, we are content in focusing on (3) by providing a technique for the evaluation of $\text{prox}_{\omega \circ B}$.

First, we observe that the minimizer \hat{y} of (3) exists and is *unique*. Indeed, this vector is characterised by the set inclusion

$$Q\hat{y} \in x - B^\top \partial\omega(B\hat{y}).$$

To make use of this observation, we introduce the affine transformation $A : \mathbb{R}^m \rightarrow \mathbb{R}^m$ defined, for fixed $x \in \mathbb{R}^d$, $\lambda > 0$, at $z \in \mathbb{R}^m$ by

$$Az := (I - \lambda BQ^{-1}B^\top)z + BQ^{-1}x$$

and the nonlinear operator $H : \mathbb{R}^m \rightarrow \mathbb{R}^m$

$$H := \left(I - \text{prox}_{\frac{\omega}{\lambda}} \right) \circ A. \quad (5)$$

The next theorem from [2] is a natural extension of an observation in [19], which only applies to the case $Q = I$.

Theorem 1. *If ω is a convex function on \mathbb{R}^m , $B \in \mathbb{R}^{m \times d}$, $x \in \mathbb{R}^d$, λ is a positive number, the operator H is defined as in (5), and \hat{y} is the minimizer of (3) then*

$$\hat{y} = Q^{-1}(x - \lambda B^\top v)$$

if and only if $v \in \mathbb{R}^m$ is a fixed point of H .

This theorem provides us with a practical tool to solve problem (3) numerically by using Picard iteration relative to the nonlinear mapping H . Under an additional hypothesis on the matrix $BQ^{-1}Q^\top$, the mapping H is non-expansive, see [2]. Therefore, Opial's Theorem [38] allows us to conclude that the Picard iterate converges to the solution of (3), see [2, 19] for a discussion of this issue. Furthermore, under additional hypotheses the mapping H is a contraction. In that case, the Picard iterate converges linearly.

We may extend the range of applicability of our observations and provide a fixed point proximal-gradient method for solving problem (1) when the regularizer has the form (2) and the error f is a *strongly smooth* convex function, that is, the gradient of f , denote by ∇f , is Lipschitz continuous with constant L . So far, the convergence of this extension has yet to be analyzed. The idea behind proximal-gradient methods, see [9, 24, 32] and references therein, is to update the current estimate of the solution x_t using the proximity operator of g and the gradient of f . This is equivalent to replacing f with its linear approximation around a point which is a function of the previous iterates of the algorithm. The simplest instance of this iterative algorithm is given in Algorithm 1. Extensions to acceleration schemes are described in [2].

Algorithm 1 Proximal-gradient & fixed point algorithm.

```

 $x_1 \leftarrow 0$ 
for  $t=1,2,\dots$  do
    Compute  $x_{t+1} \leftarrow \text{prox}_{\frac{\omega}{L}}(x_t - \frac{1}{L}\nabla f(x_t))$ 
    by the Picard process.
end for

```

2.3 Connection to the forward-backward algorithm

In this section, we consider the special case $Q = I$ and interpret the Picard iteration of H in terms of a *forward-backward algorithm* in the dual, for a discussion of the forward-backward algorithm, see for example [9]

The Picard iteration is defined as

$$v_{t+1} \leftarrow (I - \text{prox}_{\frac{\omega}{\lambda}})((I - \lambda BB^\top)v_t + Bx) \quad (6)$$

We first recall the Moreau decomposition, see, for example, [9] and references therein, which relates the proximity operators of a lower semicontinuous convex function $\varphi : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{+\infty\}$ and its conjugate,

$$I = \text{prox}_\varphi + \text{prox}_{\varphi^*} . \quad (7)$$

Using equation (7), the iterative step (6) becomes

$$v_{t+1} \leftarrow \text{prox}_{\left(\frac{\omega}{\lambda}\right)^*}(v_t - (\lambda BB^\top v_t - Bx))$$

which is a forward-backward method. We can further simplify this iteration by introducing the vector $z_t := \lambda v_t$ and obtaining the iterative algorithm

$$z_{t+1} \leftarrow \lambda \text{prox}_{\left(\frac{\omega}{\lambda}\right)^*}\left(\frac{1}{\lambda}z_t - (BB^\top z_t - Bx)\right) .$$

Using the readily verified formulas

$$\frac{1}{\lambda} \text{prox}_{\lambda g} \circ \lambda I = \text{prox}_{\frac{1}{\lambda} g \circ \lambda I}$$

and

$$\left(\frac{\omega}{\lambda}\right)^* = \frac{1}{\lambda} \omega^* \circ \lambda I$$

see, for example, [5], we obtain the equivalent forward-backward iteration

$$z_{t+1} \leftarrow \text{prox}_{\lambda \omega^*}(z_t - (\lambda BB^\top z_t - \lambda Bx)) .$$

This method is a forward-backward method of the type considered in [8, Alg. 10.3] and solves the minimization problem

$$\min \left\{ \frac{1}{2} \|B^\top z - x\|^2 + \omega^*(z) : z \in \mathbb{R}^m \right\} .$$

This minimization problem in turn can be viewed as the dual of the primal problem

$$\min \left\{ \frac{1}{2} \|u\|^2 - \langle x, u \rangle + \omega(Bu) : u \in \mathbb{R}^d \right\} \quad (8)$$

by using Fenchel's duality theorem, see, for example, [5]. Moreover, the primal and dual solutions are related through the conditions $-B^\top \hat{z} = \hat{u} - x$ and $\hat{z} \in \partial \omega(B\hat{u})$, the first of which implies that $x - \lambda B^\top \hat{v}$ equals the solution of the proximity problem (8), that is, equals $\text{prox}_{\omega \circ B}(x)$.

3 Examples of Composite Functions

In this section, we provide some examples of penalty functions which have appeared in the literature that fall within the class of linear composite functions (2).

We define for every $d \in \mathbb{N}$, $x \in \mathbb{R}^d$ and $J \subseteq \{1, \dots, d\}$, the restriction of the vector x to the index set J as $x|_J = (x_i : i \in J)$. Our first example considers the Group Lasso penalty function, which is defined as

$$\omega_{\text{GL}}(x) = \sum_{\ell=1}^k \|x|_{J_\ell}\|_2, \quad (9)$$

where J_ℓ are prescribed subsets of $\{1, \dots, d\}$ (also called the “groups”) such that $\cup_{\ell=1}^k J_\ell = \{1, \dots, d\}$. The standard Group Lasso penalty, see, for example, [36], corresponds to the case that the collection of groups $\{J_\ell : 1 \leq \ell \leq k\}$ forms a partition of the index set $\{1, \dots, d\}$, that is, the groups do not overlap. In this case, the optimization problem (4) for $\omega = \omega_{\text{GL}}$ decomposes as the sum of separate problems and the proximity operator is readily obtained by using the proximity operator of the ℓ_2 -norm to each group separately. In many cases of interest, however, the groups overlap and the proximity operator cannot be easily computed.

Note that the function (9) is of the form (2). We let $d_\ell = |J_\ell|$, $m = \sum_{\ell=1}^k d_\ell$ and define, for every $z \in \mathbb{R}^m$, $\omega(z) = \sum_{\ell=1}^k \|z_\ell\|_2$, where, for every $\ell = 1, \dots, k$ we let $z_\ell = (z_i : \sum_{j=1}^{\ell-1} d_j < i \leq \sum_{j=1}^\ell d_j)$. Moreover, we choose $B^\top = [B_1^\top, \dots, B_k^\top]$, where B_ℓ is a $d_\ell \times d$ matrix defined as

$$(B_\ell)_{ij} = \begin{cases} 1 & \text{if } j = J_\ell[i] \\ 0 & \text{otherwise} \end{cases},$$

where for every $J \subseteq \{1, \dots, d\}$ and $i \in \{1, \dots, |J|\}$, we denote by $J[i]$ the i -th largest integer in J .

The second example concerns the Fused Lasso [31], which considers the penalty function $x \mapsto g(x) = \sum_{i=1}^{d-1} |x_i - x_{i+1}|$. This function falls into the class (2). Indeed, if we choose ω to be the ℓ_1 norm and B the first order divided difference matrix

$$B = \begin{bmatrix} 1 & -1 & 0 & \dots & \dots \\ 0 & 1 & -1 & 0 & \dots \\ \vdots & \ddots & \ddots & \ddots & \ddots \end{bmatrix}$$

we get back g . The intuition behind the Fused Lasso is that it favors vectors which do not vary much across contiguous components. Further extensions of this case

may be obtained by choosing B to be the incidence matrix of a graph, leading to the penalty $\sum_{(i,j) \in E} |x_i - x_j|$. This is a setting which is relevant, for example, in online learning over graphs [13, 14].

The next example considers composition with orthogonally invariant (OI) norms. Specifically, we choose a symmetric gauge function h , that is, a norm h , which is both *absolute* and *invariant under permutations* [35] and define the function $\omega : \mathbb{R}^{d \times n} \rightarrow [0, \infty)$, at X by the formula $\omega(X) = h(\sigma(X))$, where $\sigma(X) \in [0, \infty)^r$, $r = \min(d, n)$ is the vector formed by the singular values of matrix X , in non-increasing order. An example of OI-norm are Schatten p -norms, which correspond to the case that ω is the ℓ_p -norm. The next proposition provides a formula for the proximity operator of an OI-norm. A proof can be found in [2].

Proposition 1. *With the above notation, it holds that*

$$\text{prox}_{h \circ \sigma}(X) = U \text{diag}(\text{prox}_h(\sigma(X))) V^\top$$

where $X = U \text{diag}(\sigma(X)) V^\top$ and U and V are the matrices formed by the left and right singular vectors of X , respectively.

We can compose an OI-norm with a linear transformation B , this time between two spaces of matrices, obtaining yet another subclass of penalty functions of the form (2). This setting is relevant in the context of multi-task learning. For example, in [1] h is chosen to be the *trace* or *nuclear* norm and a specific linear transformation which models task relatedness is considered. Specifically, the regulariser is given by $g(X) = \left\| \sigma \left(X \left(I - \frac{1}{n} e e^\top \right) \right) \right\|_1$, where $e \in \mathbb{R}^d$ is the vector all of whose components are equal to one.

4 Application to Support Vector Machines

In this section, we turn our attention to the important topic of support vector machines (SVMs), which are widely used in data analysis. SVMs were pioneered by the fundamental work of Vapnik [6, 10, 33] and inspired one of us to begin research in machine learning [11, 27, 26]. For that we are all very grateful to Vladimir Vapnik for his fundamental contributions to machine learning.

First, we recall the SVM primal and dual optimization problems, [33]. To simplify the presentation we only consider the linear version of SVMs. A similar treatment using feature map representations is straightforward and so will not be discussed here, although this is an important extension of practical value. Moreover, we only consider SVMs for classification, but our approach can be applied to SVM regression and other variants of SVMs which have appeared in the literature.

The optimisation problem of concern here is given by

$$\min \left\{ C \sum_{i=1}^m V(y_i w^\top x_i) + \frac{1}{2} \|w\|^2 : w \in \mathbb{R}^d \right\} \quad (10)$$

where $V(z) = \max(0, 1 - z)$, $z \in \mathbb{R}$, is the hinge loss and C is a positive parameter balancing empirical error against margin maximization. We let $x_i \in \mathbb{R}^d$, $i \in \{1, \dots, m\}$, be the input data and $y_i \in \{-1, +1\}$ be the class labels.

Problem (10) can be viewed as a proximity operator computation of the form (3), with $Q = I$, $x = 0$, $\omega(z) = C \sum_{i=1}^m V(z_i)$ and $B = [y_1 x_1 \dots y_m x_m]^\top$. The proximity operator of the hinge loss is separable across the coordinates and simple to compute. In fact, for any $\zeta \in \mathbb{R}$ and $\mu > 0$ it is given by the formula

$$\text{prox}_{\mu V}(\zeta) = \min(\zeta + \mu, \max(\zeta, 1)).$$

Hence, we can solve problem (10) by Picard iteration, namely

$$v_{t+1} \leftarrow \left(I - \text{prox}_{\frac{\omega}{\lambda}} \right) ((I - \lambda B B^\top) v_t) \quad (11)$$

with λ satisfying $0 < \lambda < \frac{2}{\lambda_{\max}(B B^\top)}$, which ensures that the nonlinear mapping is strictly contractive. Note that $v_t \in \mathbb{R}^m$ and that this iterative scheme may be interpreted as acting on the SVM dual, see Section 2.3. In fact, there is a simple relation to the support vector coefficients given by the equation $v = \frac{1}{\lambda} \alpha$. Consequently, this algorithmic approach is well suited when the sample size m is small compared to the dimensionality d . An estimate of the primal solution, if required, can be obtained by using the formula $w = -\lambda B^\top v$. Also, when $d < m$ the last equation, relating w and v , cannot be inverted. Hence, (11) is not useful in this case.

Recall that the dual problem of (10) is given [33]

$$\min \left\{ \frac{1}{2} \|B^\top \alpha\|^2 - 1^\top \alpha : \alpha \in [0, C]^m \right\}. \quad (12)$$

This problem can be seen as the computation of a generalized proximity operator of the type (3). To explain what we have in mind we use the notation \odot as the elementwise product between matrices of the same size (Schur product) and introduce the kernel matrix $K = [x_1 \dots x_m]^\top [x_1 \dots x_m]$.

Using this terminology, we conclude that problem (12) is of the form (3) with $Q = K \odot y y^\top$, $x = \mathbf{1}$ (the vector of all ones), $B = I$ and $\omega = \omega_C$, where $\omega_C(\alpha) = 0$ if $\alpha \in [0, C]^m$ and $\omega_C(\alpha) = +\infty$ otherwise. Furthermore, the proximity operator for ω is given by the projection on the set $[0, C]^m$, that is $\text{prox}_{\omega_C}(\alpha) = \min(C, \max(0, \alpha))$. These observations yield the Picard iteration

$$v_{t+1} \leftarrow \left(I - \text{prox}_{\omega_C} \right) ((I - \lambda (K^{-1} \odot y y^\top)) v_t + (K^{-1} \odot y y^\top) \mathbf{1}) \quad (13)$$

with $0 < \lambda < 2 \lambda_{\min}(K)$. This iterative scheme requires that the kernel matrix K is invertible, which is frequently the case, for example, in the case of Gaussian kernels. Another requirement is that either K^{-1} has to be precomputed or a linear system involving K has to be solved at every iteration, which limits the scalability of this scheme to very large samples. In contrast, the iteration (11) can always be applied, even when K is not invertible. In fact, when K , and equivalently $B B^\top$, is invertible

then both iterative methods (11), (13) converge linearly at a rate which depends on the condition number of K , see [2, 19].

Recall that algorithm (11) is equivalent to a forward-backward method in the dual, see Section 2.3. Thus, an accelerated variant akin to Nesterov’s optimal method and FISTA [3] could also be used. However, in the case of an invertible kernel matrix, both versions converge linearly [24] and hence it is not clear whether there is any practical advantage from the Nesterov update. Furthermore, algorithm (13) could also be modified in a similar way.

On the other hand, if $m > d$, we would directly attempt to solve the primal problem. In this case, the Nesterov smoothing method can be employed, [23]. An advantage of such a method is that it only stores $O(d)$ variables, even though it needs $O(md)$ computations per iteration. The method described above, based on Picard iteration, requires $\min(O(md), O(m^2))$ cost per iteration and stores $O(m)$ variables.

Let us finally remark that iterative methods similar to (11) or (13) can be applied to ℓ_2 regularization problems, other than SVMs, provided that the proximity operator of the corresponding loss function is available. Common choices for the loss function, other than the hinge loss, are the logistic and square loss functions leading to logistic regression and least squares regression, respectively. In particular, in these two cases, the primal objective (10) is both smooth and strongly convex and hence a linearly convergent gradient descent or accelerated gradient descent method can be used [25], regardless of the conditioning of the kernel matrix.

5 Conclusion

We presented a general approach to solve a class of nonsmooth optimization problems, whose objective function is given by the sum of a smooth term and a nonsmooth term which is obtained by linear function composition. The prototypical example covered by this setting is a linear regression regularization method, in which the smooth term is an error term and the nonsmooth term is a regularizer which favors certain desired parameter vectors. An important feature of our approach is that it can deal with a rich class of regularizers and, as shown numerically in [2], is competitive with the state of the art methods. Using these ideas, we also provided a fixed-point scheme to solve support vector machines. Although numerical experiments have yet to be done, we believe this method is simple enough to deserve attention by practitioners.

We believe that the method presented here should be thoroughly investigated both in terms of convergence analysis, where ideas presented in [34] may be valuable, and numerical performance with other methods, such as alternate direction of multipliers, see, for example, [4], block coordinate descent, alternate minimization and others. Finally, there are several other machine learning problems where ideas presented here apply. For example, in that regard we mention multiple kernel learning, see for example, [18, 28, 29, 30] and references therein, some structured sparsity

regularizers [16, 17] and multi-task learning, see, for example [1, 7, 12]. We leave these tantalizing issues for future investigation.

Acknowledgements

Part of this work was supported by EPSRC Grant EP/H027203/1, Royal Society International Joint Project Grant 2012/R2 and by the European Union Seventh Framework Programme (FP7 2007-2013) under grant agreement No. 246556.

References

1. Argyriou, A., Evgeniou, T., and Pontil, M. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.
2. Argyriou, A., Micchelli, C.A., Pontil, P., Shen, L., and Xu, Y. Efficient first order methods for linear composite regularizers, *arXiv:1104.1436*, 2011.
3. Beck, A. and Teboulle, M. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal of Imaging Sciences*, 2(1):183–202, 2009b.
4. Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning* 3(1):1–122, 2011.
5. Borwein, J. M. and Lewis, A. S. *Convex Analysis and Nonlinear Optimization: Theory and Examples*. CMS Books in Mathematics. Springer, 2005.
6. Boser, B.E., Guyon, I.M., and Vapnik, V.N. A training algorithm for optimal margin classifiers. Proc. 5th Annual ACM Workshop on Computational Learning Theory, pages 144–152, 1992.
7. Cavallanti, G., Cesa-Bianchi, N., Gentile, C. Linear algorithms for online multitask classification. *J. Machine Learning Research*, 11:2901–2934, 2010.
8. Combettes, P.L. and Pesquet, J.-C. Proximal splitting methods in signal processing. In: *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, (Bauschke, H.H. et al. Editors), pp. 185–212. Springer, 2011.
9. Combettes, P.L. and Wajs, V.R. Signal recovery by proximal forward-backward splitting. *Multiscale Modeling and Simulation*, 4(4):1168–1200, 2006.
10. Cortes, C. and Vapnik, V.N. Support-Vector Networks. *Machine Learning*, 20, 1995.
11. Evgeniou, Pontil, M., Poggio, T. *Advances in Computational Mathematics*, 13(1):1–50, 2000.
12. Evgeniou, Pontil, M., Toubia, O. A convex optimization approach to modeling heterogeneity in conjoint estimation. *Marketing Science*, 26:805–818, 2007.
13. Herbster, M. and Lever, G. Predicting the labelling of a graph via minimum p-seminorm interpolation. In *Proceedings of the 22nd Conference on Learning Theory (COLT)*, 2009.
14. Herbster, M. and Pontil, M. Prediction on a graph with the perceptron. *Advances in Neural Information Processing Systems* 19, pages 577–584, MIT Press, 2007.
15. Jenatton, R., Audibert, J.-Y., and Bach, F. Structured variable selection with sparsity-inducing norms. *arXiv:0904.3523v2*, 2009.
16. Maurer, A. and Pontil, M. Structured sparsity and generalization. *J. Machine Learning Research*, 13:671–690, 2012.
17. Micchelli, C.A., Morales, J.M., Pontil, M. A family of penalty functions for structured sparsity NIPS 2010.
18. Micchelli, C.A. and Pontil, M. Feature space perspectives for learning the kernel. *Machine Learning*, 66:297–319, 2007.

19. Micchelli, C.A., Shen, L., and Xu, Y. Proximity algorithms for image models: denoising. *Inverse Problems*, 27(4), 2011.
20. Moreau, J.J. Fonctions convexes duales et points proximaux dans un espace hilbertien. *Acad. Sci. Paris Sér. A Math.*, 255:2897–2899, 1962.
21. Mosci, S., Rosasco, L., Santoro, M., Verri, A., and Villa, S. Solving Structured Sparsity Regularization with Proximal Methods. In *Proc. European Conf. Machine Learning and Knowledge Discovery in Databases*, pp. 418–433, 2010.
22. Nesterov, Y. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Soviet Mathematics Doklady*, 27(2):372–376, 1983.
23. Nesterov, Y. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
24. Nesterov, Y. *Gradient methods for minimizing composite objective function*. CORE, 2007.
25. Nesterov, Y. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer, 2004.
26. Pontil, M., Rifkin, R.M., and Evgeniou, T. From regression to classification in support vector machines. *Proc. 7th European Symposium on Artificial Neural Networks*, pages 225–230, 1999.
27. Pontil, M. and Verri, A. Properties of support vector machines. *Neural Computation*, 10:955–974, 1998.
28. Rakotomamonjy, A. Bach, F., Canu, S., Grandvalet, Y. SimpleMKL. *J. Machine Learning Research*, 9:2491–2521, 2008.
29. Sonnenburg, S., Rätsch, G., Schäfer, C., Schölkopf, B. Large scale multiple kernel learning. *J. Machine Learning Research*, 7:1531–1565, 2006.
30. Suzuki, T. and Tomioka, R. SpicyMKL: a fast algorithm for multiple kernel learning with thousands of kernels. *Machine Learning*, 85(1):77–108, 2011.
31. Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108, 2005.
32. Tseng, P. Approximation accuracy, gradient methods, and error bound for structured convex optimization. *Mathematical Programming*, 125(2):263–295, 2010.
33. Vapnik, V. (1999). *The Nature of Statistical Learning Theory*. Springer, 1999.
34. Villa, S., Salzo, S., Baldassarre, L., Verri, A. Accelerated and inexact forward-backward splitting. *Optimization Online*, August 2011.
35. Von Neumann, J. Some matrix-inequalities and metrization of matric-space. *Mitt. Forsch.-Inst. Math. Mech. Univ. Tomsk*, 1:286–299, 1937.
36. Yuan, M. and Lin, Y. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society, Series B*, 68(1):49–67, 2006.
37. Zhao, P., Rocha, G., and Yu, B. Grouped and hierarchical model selection through composite absolute penalties. *Annals of Statistics*, 37(6A):3468–3497, 2009.
38. Zălinescu, C. *Convex Analysis in General Vector Spaces*. World Scientific, 2002.